

AUTOMATIC CONTEXT MANAGEMENT

FOR WEB APPLICATIONS WITH CLIENT SIDE CODE EXECUTION

FIELD OF THE INVENTION

5 The present invention generally relates to the field of computer applications and systems, and more specifically to a method and apparatus for testing, monitoring, and automating network applications and other fields where web client simulations are used.

BACKGROUND OF THE INVENTION

10 Testing, monitoring, automation, and other web client simulation tools, consecutively called simulation tools, often use a recorder to automatically generate scripts that simulate user activity for repeated replay. For example, load testing tools can simulate large numbers of users using multiple technologies and allow businesses to predict the behavior of their e-business web application before it is introduced for actual use, regardless of the size and complexity of the

15 application. A reliable load testing tool can simulate the activities of real users of various profiles who use web clients such as web browsers to access the web server. The load testing tool can simulate the activities of a maximum expected number of users, maintain a steady number of users over a certain period of time, or stress a single component of the web application. The load testing tool measures performance, reliability, accuracy and scalability and as such can reveal a wide

20 range of possible errors. The load testing tool simulates real users by using virtual users that effect client-server interaction on the protocol level, not the graphical user interface (GUI) level. By performing a load test, businesses can test the performance, scalability, and reliability of the web application.

A load testing tool for web applications includes a controller, multiple agents, and multiple virtual users. The controller manages and monitors the load test. The agents are connected to the controller and run on the same computer as the controller or on remote computers. Each agent hosts multiple virtual users. Each virtual user makes requests to and receives responses back from the web application on the web server. The request contains the URL of the document to be retrieved from the server, and the response contains the document itself and information about the type of document. The document can be, for example, an HTML document, an image, a PDF file, a JavaScript file, or a plain text file.

A simulation tool incorporates a replay engine which simulates virtual users with respect to the network traffic they generate, as shown in Figure 6. A script is a series of instructions which are input to the replay engine and can be in any format the replay engine understands (e.g., textual script written in a programming language, instructions stored in a database, or XML). A script can be written using a software tool like a text editor.

However, the most convenient way of generating a script is to use a recorder which generates scripts based on the HTTP(S) transactions resulting from a real user using the web application. The recorder records into one or more scripts the actions of the real user such as clicking on hyperlinks, submitting forms, transitioning back and forth in the session history (e.g. using a web browser back button), and pausing between activity.

The recorded scripts are then replayed simultaneously to simulate the user interactions of many users in order to test the server. The scripts for each virtual user are executed concurrently by the replay engine to generate the desired load on the test environment. Each component of the test environment can be monitored in real-time during the load test replay to allow the testers to view the web application's performance at any time. The results of the load test replay can be

analyzed to improve weaknesses and address errors in the web application. The scripts can be rerun to verify the adjustments.

Real users interact with a web application using a client program on a computer 101 that calls upon services provided by a server program as shown in Figure 1. The client program can be
5 a web browser that requests services from a web server 102, which in turn responds to the request by providing the services. This interaction is in the form of HTTP requests 103 and responses 104.

Virtual users do not interact with the web application using a client program since this would involve running a client program for each individual virtual user. Running separate client
10 programs for each virtual user is impractical for a load test because it wastes resources. Instead, the interactions of each virtual user with the web server take place on the protocol level. Therefore, in conventional load testing tools, the scripts display user interactions in the form of HTTP requests. The recorder records network activity between the client application and the server. Recorders used in load testing do not record activity within the client application such as
15 the movement of each user's mouse or the specific keystrokes that each user performs.

A web page, as shown in Figure 2, includes one or more documents received from the web server. Each document is received by sending one HTTP request (or more than one request when there are redirections). A web page forms a tree which has a root document 201 and leaves. A root document is a document that can have one or more embedded sub-documents 202, 203 such as
20 HTML documents, images, embedded objects, frames, scripts, applets, and style sheets. Leaves 204 are documents, such as images, style sheets, and plain text, that cannot have sub-documents. The visual representation of a web page is exactly what the real user sees when using the web browser. Documents can also contain hyperlinks 205 and forms 206, which are not separate

documents. A user can click on a hyperlink to transition to another web page having the URL associated with the hyperlink. Forms can be completed and submitted to the server by the user, thereby causing a transition to another web page having the URL associated with the form. Web pages can also include executable code that is executed by the web browser (e.g. JavaScript, VBScript, and Java applets). The executable code can be embedded in HTML documents or contained in sub-documents.

Figure 5 shows a session history, which is a sequence of web pages that are downloaded or retrieved from a client side cache. The interaction of the user with the web application is called a user session. The session begins when the user starts the browser and navigates to the web application. The session ends when the user closes the browser or navigates to a different web application. The URL of the first page 501 is specified by the user, e.g., by entering the URL in the address bar of the browser 511 or by clicking on a "bookmark". Each successive page 502, 503, 504, 505 is the result of a page transition from one page to the next page. A page transition can be caused by the user clicking a hyperlink 512, the user filling in and submitting a form 513, by the user navigating in the browser's page history by clicking the "back" or "forward" button 514, or by the browser executing client side code 515.

Although the HTTP protocol is a stateless protocol, complex web applications need to be able to relate a single request to one or more preceding requests. Such requests are common in web applications that use session IDs, shopping carts, user IDs, and other forms of state information. These forms of state information can be in the form of unique strings of characters that are placed in requests or responses, as well as the transferred documents such as HTML. The unique string appears hard-coded in each related request and response. Real clients such as web browsers can correctly identify state information in responses and can also correctly embed such

state information in subsequent requests.

For example, a session ID is a unique string of characters that a web application assigns to all responses within a period of time in which a client and a server interact. The client then returns the session ID within a request to the server so that the web application can determine which client sent the request. A shopping cart is commonly used in e-business applications to handle catalogs of items that an individual user would like to purchase. User IDs are assigned by a web application to identify a particular user. If a simulation tool cannot correctly identify a session ID, a shopping cart ID, a user ID, or other forms of state information within a response and transfer it back to the server in subsequent requests, the simulation tool does not correctly simulate real clients. This may lead to invalid test results or even errors which in fact are not errors of the web application but rather an artifact of the simulation tool being unable to simulate real clients properly.

Conventional load testing tools can only handle standardized state information called cookies. All other forms of state information can, if at all, only be handled by manually customizing the script.

Context management is the ability of a testing tool to: a) manage state information during replay by dynamically analyzing content for state information, even when the content is generated dynamically and contains state information; b) properly transfer this state information back and forth in requests and responses; c) restructure received documents (one or more HTTP(S) transactions) into web pages; and d) maintain a session history to allow virtual users to transition back and forth between web pages in the session history.

Missing or poor context management of a simulation tool can result in the mishandling of state information which ultimately can lead to a failed, unreliable, and ineffectual load test. For

example, when multiple virtual users log into the same account or session due to incorrect state management, the load on the web application does not correctly model the real intended behavior of the live application. Accurate simulation/testing tools would create one session per virtual user. Poor context management of a load testing application, in particular, can lead to inaccurate test results or failed tests. Missing automatic context management by the simulation tool must be compensated by large efforts to customize scripts manually for correct state management, thereby resulting in high costs and lost revenue.

State management may be achieved using state information which can be included as a unique string in cookies, URLs of links or embedded objects, or form fields. The string acts as a session ID or other ID such as an encryption ID or a serialized object reference to an object residing at the server which is placed in a hidden field. The string allows the server to identify the unique web session in which the original request originated or other state information, and the string is returned by the browser to the server as part of any subsequent request, thus relating the requests.

In a simulation tool without state management, the hard-coded session ID is sent to the server when replaying the script. However, since the specific session ID does not correctly identify the replayed session, this replay does not run correctly. The session ID only identifies the session ID of the recorded session and cannot be used again for the replayed session. Therefore, a script that uses such state information is unsuitable for a proper load test, and the web application will most likely generate errors during replay since sessions are usually terminated by the web server after a predetermined period of time. Load testing tools without state management generate scripts that must thus be customized, manually or via a tool, to handle state information such as session IDs in web applications in order to avoid such problems.

Conventional simulation tools use a HTTP-level replay or a page-level replay to execute scripts. A low-level, HTTP-level replay executes script instructions that specify single HTTP transactions. The information for each HTTP transaction, such as the URL and form data, is specified in the script. Because of this, session information, which is part of these URLs, and form data are hard-coded in the script and are not dynamically parsed out of documents during runtime. A HTTP-level API is therefore not suited for automatic state management and does not provide context management as well.

In contrast, a conventional simulation tool with a page-level replay executes script instructions that specify complete web pages that can include various images and frames to be downloaded. The replay engine uses a document parser (e.g. HTML parser) to obtain the URLs referencing embedded objects and frames at replay time. Since downloading a single web page using the page-level API automatically initiates HTTP requests for also downloading embedded objects or frames, the page-level script is typically shorter than the HTTP-level script. The page-level script can also automatically obtain state information which may be contained in URLs referencing frames and embedded objects in real time. Such state information is contained in URLs embedded in HTML tags of HTML documents and parsed in real time, but is not hard-coded in the page-level script .

The present invention includes a recorder which is able to identify the context of web page transitions by analyzing the HTTP(S) network traffic even for web applications that use client side code execution. Conventional recorders can also record context-full page-level scripts for web applications without client side code, but fail in many cases to do so if the web application uses client side code.

Both a web browser and the replay engine of a conventional simulation tool that is capable

of a page-level replay follow the same steps when downloading a web page from a web server. Documents are downloaded (step 401) by the client from the server as shown in Figure 4, starting with the root document. If a downloaded document is parsable (step 402) (e.g. an HTML document), a standard HTML parser parses the document (step 404). A standard HTML parser
5 parses each HTML document as it is received by the client from the web server so that embedded objects and frames can be detected and downloaded automatically. Since the conventional page-level API function call requests all embedded objects and frames in the root document, the standard HTML parser specifically looks for these embedded objects and frames in the HTML document. The standard HTML parser also looks for hyperlinks and forms in the HTML document so that
10 they can be referenced in subsequent API calls when transitioning between web pages.

As shown in Figure 3, the standard HTML parser 302 parses an HTML document 301 and can output a list of frames 311, a list of embedded documents 312, a list of hyperlinks 313, and a list of forms 314, each with their associated URLs. Each embedded object and frame is downloaded following the same procedure of Figure 4 in a recursive way. A document that is not
15 parsable stops the recursion (step 403).

The conventional recorder uses the HTML parser to analyze the HTML documents wherein a single HTML document is retrieved by one (or more when there are redirections) HTTP requests. Conventional recorders save some state information hard-coded in the API call parameters in the script..

20 Page-level replay instructions can be "context-full" or "context-less", depending on how a page transition is specified by the replay instruction. A "context-less" replay instruction can be executed by the replay engine using information only specified in the script. A context-less replay instruction includes the URL of the root document to download and optional replay instructions

that send form data without referring to a form contained in a previously downloaded web page. Both the URL and all form field names and values are specified in the script, possibly containing hard-coded state information.

The term "context-less" refers to the fact that the replay engine can execute such a replay
5 instruction without using the context of the user session executed so far. No references to previously downloaded web pages in the user session exist, and no dynamic data from previously downloaded web pages is used.

A "context-full" replay instruction is a replay instruction which refers to a previously downloaded page. The term "context-full" refers to the fact that the replay engine can execute the
10 replay instruction only within the context of the replay session up to the point where the context-full replay instruction is executed. Without a session history, the replay engine would not be able to collect the data needed to perform a context-full replay instruction.

The replay engine identifies the HTML document and all of the information associated with the particular HTML document such as the embedded images and frames that form the HTML
15 document. All of the information is downloaded in real-time during the replay. There are no session IDs hard-coded in the script.

For example, a replay instruction can download a web page by following a hyperlink contained in the previously downloaded web page. To download the web page, the replay engine uses the document parser to obtain the URL that is associated with that hyperlink. The URL
20 associated with the hyperlink is obtained in real-time during the replay. The script only contains the reference to the hyperlink, but not the URL, which might contain state information. A reference to a hyperlink may consist of several items, e.g., name of the hyperlink, ordinal number, name of containing frame, reference to the containing web page. These types of references

typically do not contain state information.

In another example, a replay instruction submits a form which is contained in a previously downloaded web page, given a reference to that form. The replay engine uses the document parser to obtain the URL associated with the form, and the names and initial values of the form fields.

5 The script only contains the name of the form and the values of form fields that are different from the original values (i.e. the values that have been edited by the user). The script does not contain the URL, which may contain state information. The script also does not contain the values of hidden or otherwise unchanged form fields, which also may contain state information. A reference to a form may consist of several items, e.g., name of the form, ordinal number, name of
10 containing frame, reference to the containing web page. These types of references typically do not contain state information.

The page-level script can eliminate many of the URLs that contain state information since usually the URL for the first HTML document is the only URL hard-coded in the script. The URL typically corresponds to the first page specified by the virtual user and typically does not
15 contain state information since it is the entry point to the web application. The other links are obtained dynamically during the replay as URLs which are obtained by parsing the downloaded HTML document during replay. These links correspond to context-full replay instructions in the script. By obtaining information dynamically, the use of hard-coded session IDs in the script can be avoided.

20 The capability of the replay engine for automatic state management by means of executing context-full replay instructions is called automatic context management. A page-level replay is suited for automatic context management for web applications that do not use client side code execution, i.e., each page transition is implemented by standard hyperlinks and form submissions.

makes it difficult to model accurately the interactions between the web application and the virtual users in a page-level API. Recorders of typical simulation tools may not be able to properly record context-full functions when code is executed on the client side and will record context-less functions instead.

5 The recorder cannot determine the context of an API function when it observes an HTTP request that cannot be interpreted as an embedded object, a frame, a link, a form, or any other entity that the standard HTML parser can detect. The standard HTML parser does not have the ability to identify the URLs generated during client side code execution. Therefore, the recorder records context-less function calls corresponding to the non-interpreted HTTP request.

10 The standard HTML parser cannot recognize a URL that is generated by executing client side code and therefore cannot associate the URL of the recorded HTTP transaction with any URL in the session history. Therefore, the recorder records a hard-coded URL in the API function call.

 The scripts that are generated by a recorder using a standard HTML parser produce an unreliable load test replay when the replay engine executes a context-less replay instruction using
15 hard-coded state information. The standard HTML parser is only able to identify URLs found in HTML tags such as hyperlinks and references to frames and embedded objects.

 Client-side code causes actions that are performed by the client, rather than by executing web application code on the server. Code executed on the client side gives web developers increased functionality over standard HTML techniques and provides the client side portion of the
20 web application with abilities such as loading a web document, loading embedded objects, and modifying HTML code and form values.

 Conventional load testing tools exhibit poor context management when the conventional page-level recorder generates context-less function calls after code is executed on the client side.

less API function calls.

As shown in Figure 13, the recorder processes a HTTP transaction by inspecting the transaction (step 1301) using the session history 1302 to identify the role of the HTTP transaction within the session history.

5 If the HTTP transaction corresponds to an embedded object (step 1311), no replay instruction is added to the script. The session history is updated to reflect the fact that this embedded object has been downloaded.

10 If the HTTP transaction corresponds to a frame (step 1312), no replay instruction is added to the script. The session history is updated to reflect the fact that this frame has been downloaded.

15 If the HTTP transaction corresponds to a hyperlink (step 1313), a new context-full script instruction "FollowHyperlink" is recorded in the script along with parameters that allow the replay engine to reference the hyperlink during replay. A new web page is added to the session history resulting from following the hyperlink. However, this new web page in the session history is incomplete and will be populated with embedded documents and frames as the recorder processes the upcoming HTTP transactions.

20 If the HTTP transaction corresponds to a form submission of an existing form (step 1314), a new context-full script instruction "SubmitForm" is recorded in the script along with parameters that allow the replay engine to reference the form during replay and the names and values of the form fields that have been edited by the user. A new web page is added to the session history resulting from the form submission. However, this new web page in the session history is incomplete and will be populated with embedded documents and frames as the recorder processes the upcoming HTTP transactions.

If the HTTP transaction corresponds to neither a hyperlink nor a form submission, but contains form data (step 1315), a new context-less script instruction "SendForm" is recorded in the script along with the complete form data that is to be used for sending the form during script replay, including the URL to use and the complete list of names and values of the form fields. A
5 new web page is added to the session history resulting from the form submission. However, this new web page in the session history is incomplete and will be populated with embedded documents and frames as the recorder processes the upcoming HTTP transactions.

If the HTTP transaction corresponds to neither a hyperlink nor a form submission and does not contain form data (step 1316), a new context-less script instruction "LoadPage" is recorded in
10 the script, along with the URL that is to be used for loading the page during script replay. A new web page is added to the session history representing the new web page. However, this new web page in the session history is incomplete and will be populated with embedded documents and frames as the recorder processes the upcoming HTTP transactions.

If the HTTP transaction corresponds to neither a hyperlink nor a form submission (steps
15 1315 and 1316), the recorder records context-less script instructions to the script. The way the conventional recorder handles these cases is the reason why hard-coded state information is incorporated in recorded scripts.

In a conventional recorder, each document in the session history is inspected to find forms that exactly match the form being submitted in order to find a form that can be used for a context-
20 full SubmitForm replay instruction. Forms match exactly if the action URL (i.e., the URL that defines where to send the data in the submitted form when the submit button is clicked or a similar action is performed) is identical, and the form being submitted contains all form fields of the form from the document in the session history.

However, the form being submitted may contain additional form fields not present in the form in the document in the session history. For instance, web browsers implicitly add the form fields "x" and "y" which contain the coordinates of the mouse click. Additionally, the form field values in the form being submitted may be different from the form field values in the form in the document in the session history because the user may have edited form field values, but the forms are considered identical since the form being submitted contains all form fields of the form from the document in the session history.

Only if an identical form is found, will the conventional recorder be able to record a context-full replay instruction such as "SubmitForm" with a reference to the form from the session history. If no such form is found in the session history, a conventional recorder records a context-less replay instruction such as "SendForm", which also requires recording in the script a complete specification of the form without using any dynamic information.

Scripts can be customized by the tester after recording the script, manually or by using a software tool. However, this method of context management is complex, error-prone, and time-consuming and wastes quality assurance (QA) resources.

SUMMARY OF THE INVENTION

The present invention relates to a method and apparatus for providing automatic context management for simulating real users with virtual users for testing and monitoring web applications, including those web applications that execute code on the client side, without requiring the actual execution of client side web application code or the execution of the client within the testing, monitoring, or simulation tool. Simulation tools with automatic context management according to the present invention can record and replay context-full scripts that do

not require manual customization and are capable of handling state information even for web applications that execute code on the client side. These scripts are able to realistically mimic complex web application sessions on the network HTTP layer.

The present invention includes a context-full page-level replay for a simulation tool that
5 uses a scripting language to model end-user behavior causing network activity between a web application and a user's web browser (client), an extensible document parser, a recorder that automatically records context-full scripts, and a context-full replay engine which is able to execute the context-full page-level API calls.

An enhanced, extensible document parser can parse URLs that standard HTML parsers
10 would overlook. The extensible document parser searches the entire HTML document for standard and nonstandard embedded objects, hyperlinks, and forms. Nonstandard hyperlinks and forms can be embedded, for example, in JavaScript code in HTML documents or in Java applets.

The extensible document parser can locate URLs in a client side code. These URLs are detected during the script record process and recorded into the script without hard-coded state
15 information. During the replay process, the otherwise hard-coded state information is obtained in real-time.

During the record process, if the recorder cannot associate a URL with a hyperlink, frame or form in the session history, the extensible document parser determines the appropriate parser extensions to be recorded into the script, so that the URL can be obtained dynamically during
20 script replay.

For example, the parser extension can include a name for the link, a left boundary string, and a right boundary string. The left and right boundary strings specify the unique strings of characters that appear to the left and right of the URL, respectively. Using the left and right

boundary strings, the extensible document parser can search for and detect the URL in any document. Whenever the left and right boundary strings appear in an HTML document, the extensible document parser identifies that the parsed URL located between the two strings corresponds to the link having the name specified in the parser extension.

5 The recorder records parser extensions into the script for use by the replay engine. The recorder records the name and boundary value fields in a parser extension in the script in an API function call so that the replay engine can parse the URL that otherwise could not be parsed. Each set of parser extensions is specific to a particular parsed URL and avoids the use of hard-coded URLs in the recorded script.

10 The extensible document parser parses for nonstandard embedded objects, hyperlinks, and forms to find URLs and forms that may contain session IDs and other forms of state information that a recorder using a standard HTML parser would include in the script.

 The present invention includes form merging instructions for the replay engine. The form merging instructions are a method of describing how a form contained in a previously downloaded
15 web page must be merged with a form defined in the script to obtain the form that must be submitted.

 In the presence of client side code execution, the user may submit a form that does not correspond exactly to any form in the HTML web pages in the user session at that time. In addition to changing all HTML in a web page, JavaScript can modify HTML forms prior to
20 submission by modifying form field values, modifying hidden fields, adding or removing form fields, renaming form fields, or changing the action URL.

 The incorporation of code executed on the client side in the form submission can result in an API function that includes hard-coded state information such as session IDs in the script.

The present invention also includes a method of using a recorder to record scripts that use both a parser addition and form merging instructions so that the recorded result is a script capable of automatic context management, but does not contain hard-coded state information. The recorder can automatically choose parser additions, along with configuration parameters for the parser additions, from a library of available additions. The recorder chooses those additions that are needed for the web application being recorded.

The recorder can also automatically detect the form in the user session that is most similar to a form being submitted, and generate form merging instructions into the script so that the replay engine will submit the correct form when executing the script. This recording process is called fuzzy form detection.

The replay engine follows the form merging instructions on how to merge the two forms to produce a form submission. The recorder generates these form merging instructions into the recorded script. The instructions specify which forms, form definitions, form fields, and alternate action URLs to use. The form is modeled using context-full form functions for the form fields that incorporate state information.

The recorder does not need any pre-configuration before recording, no user intervention during recording, and the scripts do not need customization by the user after recording.

The recorder automatically records the scripts without the need for user customization or a separate configuration. No manual pre-configuration for specific handling techniques is necessary since the recorder can automatically adapt to formerly unknown situations. There is no need to manually add parsing instructions for state information to the script because there is no state information statically incorporated in the recorded script.

Automatic context management using an extensible document parser and fuzzy form

detection and merging enables the automatic use of dynamic data in a script. No user configuration or customization is necessary before or after recording.

Automatic context management improves the ease-of-use and overall quality of a simulation tool and improves the productivity of the QA process.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a block diagram of the interaction between a client and a server;

Figure 2 is a block diagram of the structure of a web page;

10 Figure 3 is a diagram of a conventional document parser;

Figure 4 is a flowchart of a method for downloading a web page from a web server;

Figure 5 is a diagram of a sample user session;

Figure 6 is a diagram of the record/replay process for a simulation tool;

Figure 7 is a diagram of an extensible document parser of the present invention;

15 Figure 8 is a diagram of a method for merging forms according to the present invention;

Figure 9 is a diagram of input to a recorder according to the present invention;

Figure 10 is a diagram of output from a recorder according to the present invention;

Figure 11 is a diagram of a recorder of the present invention;

Figure 12 is a flowchart of a method for operating a recorder according to the present
20 invention;

Figure 13 is a flowchart of a method for processing HTTP(S) transactions using a conventional recorder; and

Figure 14 is a diagram of interfaces of a parser addition according to the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

EXTENSIBLE DOCUMENT PARSER

The present invention includes an extensible document parser 702 as shown in Figure 7, which can parse URLs by making use of the parser addition 1401 that standard document parsers overlook. Document parsers parse documents according to the syntax of the document type. The extensible document parser 702 has added functionality over the standard document parser since the extensible document parser 702 has a plug-in interface to include one or more parser additions 703.

The parser addition 703 is a DLL (dynamic link library) having an interface that can be plugged into the extensible document parser 702. The parser addition 703 is used by the extensible document parser 702 to parse additional frames 721, embedded objects 722, hyperlinks 723 and forms 724 from the document in addition to the standard frames 711, embedded objects 712, hyperlinks 713, and forms 714 from the document.

The extensible document parser 702 passes the document to parse 701 and additional input parameters to the parser addition 703. The parser addition 703 reports the parsing results back to the extensible document parser 702 by sending output parameters through the interface in the parser addition 703 that is plugged into the extensible document parser 702.

The parser addition has another plug-in interface (1403, Figure 14) used by the recorder. By using the parser addition, the recorder can record a script that contains parser extensions. A parser extension is a replay instruction that specifies parser additions and configuration data for the parser additions. The recorder, by invoking this plug-in interface 1403 of the parser addition, determines the appropriate parser extensions to record into the script, so that a URL or form can be

parsed during replay. Parser extensions can be in effect during the entire script execution or for individual page-level replay instructions only.

The parser extension specifies a parser addition and configuration data for this parser addition. In one embodiment, the configuration data for a parser addition consists of a left and a right boundary string. The parser addition first searches for the given left boundary string within the document. At each occurrence of the left boundary string, the parser addition searches for the next occurrence of the right boundary string. The text fragment found between the boundary strings is the resulting parsed URL.

In this embodiment, the configuration data for the parsing addition are the left and right boundary strings. Other types of parser additions (e.g. a parser addition that parses the document according to a regular expression) might have different kinds of input parameters (e.g. a regular expression).

PARSER ADDITION

A parser addition 1401 shown in Figure 14 is a software module that provides two interfaces 1402, 1403: one for use by the extensible document parser and one for use by the recorder, respectively. The preferred embodiment is a dynamic link library (DLL) with the interfaces being exported functions. However, any technology capable of implementing plug-ins is suitable, e.g., Microsoft COM or an object-oriented class framework.

The interface of the parser addition 1401 for the extensible document parser 1402, referred to as "ParseDocument", receives the document to parse and additional configuration parameters that completely depend on the parser addition. The parameters for a parser addition can be anything that is meaningful to parameterize the particular algorithm that the parser addition

performs. The parser addition (1401, Figure 14; 703, Figure 7) outputs a list of embedded objects 721, frames 722, hyperlinks 723, and forms 724.

The interface of the parser addition 1401 for the recorder 1403, referred to as "DetectParameters", is used by the recorder. The input to the parser addition 1401 through this
5 interface is a document and either a URL or a form. The output from the parser addition 1401 through this interface includes: a) parameters that can be passed to the interface ParseDocument along with the given document so that this interface can parse the given URL or form; or b) a notification that no suitable parameters can be detected.

The recorder according to the present invention chooses a parser addition from the library
10 of available parser additions and parameters for the chosen parser addition so that the URL or form that would otherwise be hard-coded in the script can be parsed during script replay.

The recorder relies on the parser additions to detect suitable parameters. The recorder queries each one of the parser additions available in the library of parser additions 1112 using the interface DetectParameters 1403 of each parser addition.

15 For each parser addition in the library and for each document in the session history, the recorder queries the parser addition if there are suitable parameters for the parser addition so that the parser addition can parse the required URL or form from the given document. Some parser additions will succeed and provide a set of suitable parameters for some of these documents.

The result is a set of triplets, each one including a document, a parser addition, and
20 parameters for the parser addition. Among this set, the recorder chooses one triplet which is the most suitable to be used. This choice may depend on several criteria related to both the document and the parser addition. For example, a document contained in the most recently accessed web page in the session history is preferred to a document contained in a web page farther back in the

session history. Also, parser additions can be tagged with an estimation of the resource consumption (CPU, memory) the parser addition will need during script replay, and the recorder chooses the parser addition with the lowest estimated resource consumption.

It might still be possible that no parser addition is able to report success. In this case, the recorder of the present invention records a context-less replay instruction in the script, like a conventional recorder. However, the probability of the recorder of the present invention recording a context-less replay instruction decreases significantly as more parser additions with different parsing algorithms are available.

If the recorder can successfully choose parser additions and parameters for the parser additions, the recorder records a replay instruction to the script that specifies that the chosen parser addition with the chosen parameters is to be used during script replay and during the execution of the replay instruction that downloads the web page containing the document in which the parser addition can parse the required URL or form. Such a replay instruction is called "parser extension".

The recorder then records a context-full replay instruction such as "FollowHyperlink" or "SubmitForm" instead of the context-less replay instruction such as "LoadPage" or "SendForm". In other words, the recorder avoids context-less script instructions (1315, 1316, Figure 13) and instead records context-full script instructions (1313, 1314, Figure 13).

For example, an embodiment of a parser addition is a boundary searching parser addition. The interface "ParseDocument" of the boundary searching parser addition receives a left and a right boundary string in addition to the document to parse. The parser addition first searches for the given left boundary string within the document. At each occurrence of the left boundary string the parser addition searches for the next occurrence of the right boundary string. The text

script need to also specify a different action URL. The recorder finds a suitable parser addition and parameters for the parser addition so that a replay instruction can be recorded in the script to dynamically parse the different action URL during replay. The recorder then records form merging instructions containing a reference to the parsed action URL instead of the hard-coded action URL.

FORM MERGING

The present invention includes the process of form merging. Form merging instructions are contained in the script and used by the replay engine to determine how a form contained in a previously downloaded web page is merged with a form defined in the script to obtain the form that must be submitted to the web application.

The replay engine performs form merging 804 to obtain a form to be submitted 803. The form to be submitted 803 is constructed by merging a document form 801 and a script form 802 according to form merging instructions, which are part of the script form 802. The document form 801 is a form obtained by the extensible document parser and contained in a web page previously downloaded during the replay of the script. The script form 802 is a form defined in the script. While merging the document form 801 and the script form 802, the replay engine follows form merging instructions recorded in the script to construct the form to be submitted.

The form merging instructions in the script can be a combination of, but are not limited to, the following:

- Use the form field value specified in the form obtained by the document parser (i.e., the dynamic document value) for a form field of the same name in the form to be submitted.

field "quantity" 822 from the form to be submitted. The form merging instruction 802 "use script value" specified using the fuzzy form detection process produces the result wherein the submitted form uses the scripted form field value "2" for form field "quantity" 823 from the script.

5 Form field "property" 831 from the previously downloaded web page is compared to form field "property" 832 from the form to be submitted. The form merging instruction 802 "remove form field" specified using the fuzzy form detection process produces the result wherein the submitted form does not contain the form field "property".

10 The missing form field "color" from the previously downloaded web page is compared to form field "color" 842 from the form to be submitted. The form merging instruction 802 "add form field, use script value" specified using the fuzzy form detection process produces the result wherein the submitted form contains the form field "color" 843 and uses the form field value "blue" from the script.

15 Form field "SID" 851 from the previously downloaded web page is compared to form field "SID" 852 from the form to be submitted. The form merging instruction 802 "use document value" specified using the fuzzy form detection process produces the result wherein the submitted form uses the dynamic form field value for "SID" 853 from the document.

20 The missing form field "action" from the previously downloaded web page is compared to form field "action" 862 from the form to be submitted. The form merging instruction 802 "add form field, use script value" specified using the fuzzy form detection process produces the result wherein the submitted form contains the form field "action" 863 and uses the form field value "add to basket" from the script.

With a replay engine capable of form merging, it is possible to model web page transitions by a context-full replay instruction in situations where this would not be possible with a prior

replay engine.

Examples for such situations are web application which embed incomplete forms within HTML documents, and modify these forms by means of client side code execution, e.g. JavaScript, prior to submitting such a form. Form modifications performed by client side code may be: adding form fields, removing form fields, renaming form fields or changing values of form fields.

RECORDER

The present invention includes a recorder that can choose suitable parser additions from a library of available parser additions that are needed for generating context-full replay instructions. The recorder of the present invention also performs fuzzy form detection.

The input of the recorder 1101 is a sequence of HTTP transactions as shown in Figure 9. Each HTTP transaction is either a simple HTTP transaction 901 or a compound HTTP transaction 902.

A simple HTTP transaction includes meta data 910 (e.g., timestamps and IP addresses), a HTTP request header 911, a HTTP request body 912, a HTTP response header 913, and a HTTP response body 914. The HTTP response body 914 is the document that has been downloaded by the simple HTTP transaction.

A compound HTTP transaction includes a list of simple HTTP transactions that are related by HTTP redirections (HTTP status codes "301", "302") and/or authentications (HTTP status codes "401", "407").

The output of the recorder is a script 1001 consisting of a sequence of replay instructions 1010, 1011, 1012, 1013, 1014, 1099 as shown in Figure 10. The replay instructions may contain

